

Lognormal Distributions and Geometric Averages of Positive Definite Matrices

Armin Schwartzman

Department of Statistics, North Carolina State University

July 29, 2014

Abstract

This article gives a formal definition of a lognormal family of probability distributions on the set of symmetric positive definite (PD) matrices, seen as a matrix-variate extension of the univariate lognormal family of distributions. Two forms of this distribution are obtained as the large sample limiting distribution via the central limit theorem of two types of geometric averages of i.i.d. PD matrices: the log-Euclidean average and the canonical geometric average. These averages correspond to two different geometries imposed on the set of PD matrices. The limiting distributions of these averages are used to provide large-sample confidence regions for the corresponding population means. The methods are illustrated on a voxelwise analysis of diffusion tensor imaging data, permitting a comparison between the various average types from the point of view of their sampling variability.

1 Introduction

A scalar random variable is called lognormal if its logarithm is normally distributed. Analogously, a random positive definite (PD) matrix may be called lognormal if its matrix logarithm, a random symmetric matrix, has a matrix-variate normal distribution. The resulting distribution on the set of PD matrices can be naturally called the PD-matrix-variate lognormal distribution.

While useful as a probability model for PD matrix data (Schwartzman, 2006), it is argued here that this distribution plays a more central role in multivariate statistics. Just like the scalar lognormal distribution may be obtained as the limiting distribution of the geometric average of i.i.d. random variables (Johnson and Kotz, 1970), it is shown here that two types of PD-matrix-variate lognormal distributions arise as the limiting distributions of two types of geometric averages of i.i.d. random PD matrices. The main goal of this paper is to justify and make explicit the definition of such distributions. Moreover, it is shown how these distributions can be used to construct confidence regions (CR's) for those geometric averages in data analysis.

The first geometric average considered in this paper, called log-Euclidean average, is obtained by taking the average of the PD matrices after a matrix-log transformation, and transforming the result back to the original space (Arsigny et al., 2006; Schwartzman, 2006). Application of the multivariate central limit theorem (CLT) in the log domain leads directly to the definition of a PD-matrix-variate lognormal distribution, which we call Type I, so that the log-Euclidean average can be regarded as approximately lognormal for large sample sizes. The second geometric average, called here canonical geometric average, is the Fréchet average associated with the so-called canonical, affine-invariant or generalized Frobenius Riemannian

metric defined on the set of PD matrices (Kobayashi and Nomizu, 1996; Lang, 1999; Moakher, 2005; Smith, 2005; Schwartzman, 2006; Lenglet et al., 2006; Fletcher and Joshi, 2007; Zhu et al., 2009). This case requires application of a Riemannian manifold version of the CLT, and the limiting distribution leads to the definition of a second type of lognormal distribution, called here Type II. The Type I and Type II distributions coincide when the mean parameter in both is a multiple of the identity matrix.

As a side benefit, the theoretical analysis above allows establishing some properties of these geometric averages and their corresponding geometric population means, such as uniqueness, consistency and invariance. The two geometric averages are contrasted with the simpler arithmetic or Euclidean average of PD matrices, which is also consistent but asymptotically normal. Their different properties are described in terms of the geometry of the set of PD matrices.

In classical multivariate statistics, the most common probability model for PD matrices has been the Wishart distribution. Related parametric models for PD matrices include the multivariate gamma, inverse Wishart, and inverse multivariate gamma, as well as noncentral versions of these distributions (Anderson, 2003; Gupta and Nagar, 2000; Muirhead, 1982). These models assume that the observations are real-valued multivariate vectors and have as a goal to model their population or sample covariance matrix. The idea of using a matrix logarithm of PD matrices to flexibly model covariance matrices in a similar multivariate context was explored by Leonard and Hsu (1992) and Chiu et al. (1996).

Today, new data exist where the observations themselves (after pre-processing) take the form of PD matrices. In Diffusion Tensor Imaging (DTI), a special modality of magnetic resonance imaging, each voxel contains a random 3×3 PD matrix called diffusion tensor (Basser and Pierpaoli, 1996; LeBihan et al., 2001; Bammer et al., 2002). In the brain, the diffusion tensor describes the local spatial pattern of water diffusion, its eigenvalues related to the type of tissue and its health, and its eigenvectors related to the underlying spatial direction of neural fibers. In tensor-based morphometry (TBM), brain images are spatially registered to a template and the transformation is summarized at each voxel by a 3×3 PD matrix called deformation tensor (Joshi et al., 2007; Leporé et al., 2008).

Both the log-Euclidean and canonical geometric averages of PD matrices, as well as the simpler Euclidean average, have been used for DTI data (Arsigny et al., 2006; Schwartzman, 2006; Fletcher and Joshi, 2007; Zhu et al., 2009; Schwartzman et al., 2010) and TBM data (Leporé et al., 2008). The results in this paper enable the construction of CR's for these averages. Interesting questions arise about how to visualize these 6-dimensional regions, where the elements of these regions and the averages themselves are 6-dimensional objects represented as 3-dimensional ellipsoids. In order to capture the largest variation inside the CR, here we consider the display of the two ellipsoids at the boundary of the CR along its first principal axis. Interestingly, finding such boundary points for the canonical geometric average requires solving a special case of the continuous-time algebraic Ricatti equation, whose solution is in itself a form of canonical geometric average.

As a specific data example, the three types of average are compared in an analysis of DTI images corresponding to a group of 34 10-year-old children (Dougherty et al., 2007; Schwartzman et al., 2010). The analysis shows that the log-Euclidean and canonical geometric averages are statistically indistinguishable, whereas the Euclidean and log-Euclidean are within their respective ranges of sampling variability for this sample size for most voxels. These results suggest that the theoretical debate about the choice of average (Whitcher et al., 2007; Dryden et al., 2009; Pasternak et al., 2010) may sometimes be of no practical importance, making the Euclidean or log-Euclidean preferable for their computational simplicity.

The parametric approach followed in this paper contrasts and complements the nonpara-

metric approach of Osborne et al. (2013); Ellingson et al. (2013), which uses bootstrapping to establish significance in two-sample tests of canonical geometric means but is restricted to analysis of a single voxel because of computational complexity. In contrast, the parametric approach followed here can be easily applied in parallel to thousands of voxels in the entire brain volume.

The rest of the paper is organized as follows. Section 2 introduces the Euclidean and log-Euclidean averages, as well as the lognormal distribution of Type I. Section 3 presents the canonical geometric average and the resulting lognormal distribution of Type II. Section 4 presents the DTI example, concluding with a discussion in Section 5. Relevant auxiliary material appear in Appendix A and proofs are placed in Appendix B.

2 The log-Euclidean average and the PD-matrix lognormal distribution of Type I

2.1 The set of PD matrices

It is useful to think of the set $\text{Sym}^+(p)$ of $p \times p$ real PD matrices in relation to the set $\text{Sym}(p)$ of $p \times p$ real symmetric matrices in three different ways. First, $\text{Sym}^+(p)$ is a simply connected open subset of $\text{Sym}(p)$, whose boundary is the set of $p \times p$ symmetric matrices with one or more zero eigenvalues. It is also a convex cone in $\text{Sym}(p)$ because $X_1, X_2 \in \text{Sym}^+(p)$ implies $a_1 X_1 + a_2 X_2 \in \text{Sym}^+(p)$ for any scalars $a_1, a_2 > 0$. However, it is not a vector space; e.g. the negative of a PD matrix and some linear combinations of PD matrices are not PD. Figure 1 shows the shape of the cone for $p = 2$ by embedding $\text{Sym}^+(2)$ in \mathbb{R}^3 .

Second, $\text{Sym}^+(p)$ and $\text{Sym}(p)$ have a one-to-one correspondence via the matrix exponential and logarithm

$$\text{Sym}^+(p) = \exp[\text{Sym}(p)], \quad \text{Sym}(p) = \log[\text{Sym}^+(p)], \quad (1)$$

(see Appendix A.2). Third, $\text{Sym}^+(p)$ is a differentiable manifold whose tangent space at any point can be identified as a copy of $\text{Sym}(p)$. Furthermore, it can be given a Riemannian manifold structure by defining a metric, which in this paper will be Euclidean, log-Euclidean, or canonical (affine-invariant) (see Appendix A.3 and A.4). These relationships will play various roles in the development below.

A useful geometrical representation of the elements of $\text{Sym}^+(p)$ themselves is the association of $X \in \text{Sym}^+(p)$ with the ellipsoid in \mathbb{R}^p

$$\mathcal{E}(X) = \{w \in \mathbb{R}^p : w' X^{-1} w = 1\}. \quad (2)$$

This ellipsoid can be thought of as an isosurface of a multivariate normal distribution in \mathbb{R}^p with covariance X . This representation is useful for visualization of X as the semiaxes of $\mathcal{E}(X)$ correspond to the eigenvectors of X scaled by the square root of the eigenvalues. The representation of some elements of $\text{Sym}(2)$ as ellipses is illustrated in Figures 1 and 2.

2.2 The Euclidean average of PD matrices and the symmetric-matrix normal distribution

Before presenting the log-Euclidean average, it is helpful to first consider the Euclidean average for comparison. Given PD matrices $X_1, \dots, X_n \in \text{Sym}^+(p)$, the convexity of $\text{Sym}^+(p)$ implies




Figure 1: A 2×2 matrix X with diagonals a, b and off-diagonal c is PD if and only if $a > 0$, $b > 0$ and $ab - c^2 > 0$. The set of valid triplets $\text{vecd}(X) = (a, b, \sqrt{2}c)'$ (see also Appendix A.1) is an open cone in \mathbb{R}^3 . For example, the diagonal matrices $X_1 = \text{Diag}(0.9, 0.1)$ and $X_2 = \text{Diag}(0.1, 0.9)$ are represented by the triplets $\text{vecd}(X_1) = (0.9, 0.1, 0)'$ and $\text{vecd}(X_2) = (0.1, 0.9, 0)'$ respectively. Because this set is convex, interpolation between any two points is permitted, and the Euclidean average $\bar{X} = \text{Diag}(0.5, 0.5)$ is PD. Extrapolation, however, might result in matrices that are not PD: the green straight line connecting X_1 and X_2 extends beyond the boundaries of the cone. The yellow hyperbola connecting X_1 and X_2 is the matrix exponential of the straight line between $\log X_1$ and $\log X_2$, always inside the cone by construction. The midpoint along this line is the log-Euclidean average (8) equal to $\hat{X} = \text{Diag}(0.3, 0.3)$. In this case, it also coincides with the canonical geometric average defined via (21).

that the Euclidean average

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i \quad (3)$$

is also an element of $\text{Sym}^+(p)$ (Figure 1). Suppose $X_1, \dots, X_n \in \text{Sym}^+(p)$ are pairwise i.i.d. such that $E(X_i) = M \in \text{Sym}^+(p)$ (capital μ) and $\text{Cov}(\text{vecd}(X_i)) = \Sigma \in \text{Sym}^+(q)$. Here, $\text{vecd}(\cdot)$ can be in general any vectorization operator, but we prefer the specific operator defined by (29) in Appendix A.1. Then by the multivariate CLT, as $n \rightarrow \infty$,

$$\bar{X} \rightarrow M, \quad \sqrt{n} \text{vecd}(\bar{X} - M) \Rightarrow N(0, \Sigma) \quad (4)$$

in probability and in distribution, respectively. Thus $\bar{X} \in \text{Sym}^+(p)$ is a consistent estimator of M and, for large n , is approximately normally distributed in $\text{Sym}(p)$ with mean M

and covariance Σ/n (Σ can be estimated consistently by the sample covariance matrix of $\text{vecd}(X_1), \dots, \text{vecd}(X_n)$ if finite fourth order moments are also assumed).

In this first view, the PD matrices $X_1, \dots, X_n \in \text{Sym}^+(p)$ are treated as if they were mere symmetric matrices. As n increases, the distribution of \bar{X} becomes more concentrated around M and the PD constraints at the boundary of $\text{Sym}^+(p)$ become asymptotically irrelevant. However, since the normal distribution is not defined on $\text{Sym}^+(p)$ for any n because of the boundary constraints, the asymptotic normal distribution of \bar{X} is technically defined on the tangent space to $\text{Sym}^+(p)$ at M , equal to $\text{Sym}(p)$. With some abuse of notation, we write this as

$$\bar{X} \sim N(M, \Sigma/n), \quad (5)$$

where we use the notational shorthand

$$X \sim N(M, \Sigma) \quad \Leftrightarrow \quad \text{vecd}(X) \sim N(\text{vecd}(M), \Sigma). \quad (6)$$

The specific choice of the vectorization operator $\text{vecd}(\cdot)$ given by (29) gives an elegant definition of the symmetric-matrix-variate normal distribution above (Schwartzman et al., 2008), which in the case of identity covariance reduces to the random matrix model called Gaussian Orthogonal Ensemble (Mehta, 1991).

Based on (4), a confidence region (CR) for M of asymptotic level α is given by the q -dimensional ellipsoid

$$\text{CR}_{n,\alpha}(\bar{X}) = \left\{ M \in \text{Sym}^+(p) : n \text{vecd}(\bar{X} - M)' \hat{\Sigma}^{-1} \text{vecd}(\bar{X} - M) \leq \chi_{q,1-\alpha}^2 \right\}, \quad (7)$$

where $\chi_{q,1-\alpha}^2$ is the $1 - \alpha$ quantile of the χ^2 distribution with q degrees of freedom and $\hat{\Sigma}$ is a consistent estimator of Σ such as the sample covariance matrix of $\text{vecd}(X_1), \dots, \text{vecd}(X_n)$.

2.3 The log-Euclidean average of PD matrices

A second view of $\text{Sym}^+(p)$ is as the image of $\text{Sym}(p)$ via the matrix exponential. Just like positive scalar data is sometimes log-transformed before analysis, the analogous approach to analysis of PD matrices is to transform them into symmetric matrices by a matrix log transformation (Arsigny et al., 2006; Schwartzman, 2006; Fletcher and Joshi, 2007; Zhu et al., 2009; Schwartzman et al., 2010; Leporé et al., 2008). Given PD matrices $X_1, \dots, X_n \in \text{Sym}^+(p)$, mapping them to $\text{Sym}(p)$ via the matrix logarithm, computing their Euclidean average there and then mapping the result back to $\text{Sym}^+(p)$ results in the log-Euclidean average

$$\hat{X} = \exp \left(\frac{1}{n} \sum_{i=1}^n \log X_i \right). \quad (8)$$

The symbol \hat{X} reflects that this is a “curved” version of the Euclidean average (Figure 1).

Given a random PD matrix X with probability measure Q on $\text{Sym}^+(p)$, we can define the log-Euclidean mean

$$M = \exp [E(\log X)], \quad (9)$$

where the expectation is taken with respect to the measure Q , so that the log-Euclidean average (8) corresponds to the empirical measure $\hat{Q}_n = (1/n) \sum_{i=1}^n \delta(X_i)$. A direct consequence of the analytical property (33) in Appendix A.2 is that the log-Euclidean mean and the log-Euclidean average are equivariant under similarity transformations and matrix inversion. Let $GL(p)$ denote the set of $p \times p$ real invertible matrices (general linear group). If X is a random PD matrix with log-Euclidean mean M , then:

1. The log-Euclidean mean of AXA^{-1} is AMA^{-1} , for $A \in GL(p)$.
2. The log-Euclidean mean of X^{-1} is M^{-1} .

Similarly, given PD matrices X_1, \dots, X_n with log-Euclidean average (8):

1. The log-Euclidean average of $AX_1A^{-1}, \dots, AX_nA^{-1}$ is $A\widehat{X}A^{-1}$, for $A \in GL(p)$.
2. The log-Euclidean average of $X_1^{-1}, \dots, X_n^{-1}$ is \widehat{X}^{-1} .

2.4 The log-Euclidean average and the PD-matrix lognormal distribution of Type I

Suppose now that $X_1, \dots, X_n \in \text{Sym}^+(p)$ are pairwise i.i.d. PD matrices such that $E(\log X_i) = \log M \in \text{Sym}(p)$ and $\text{Cov}(\text{vecd}(\log X_i)) = \Sigma \in \text{Sym}^+(q)$. Let \widehat{X} denote their log-Euclidean average, given by (8). Then, applying the multivariate CLT to the transformed matrices $\log X_1, \dots, \log X_n$ gives that, as $n \rightarrow \infty$,

$$\widehat{X} \rightarrow M, \quad \sqrt{n} \text{vecd}(\log \widehat{X} - \log M) \Rightarrow N(0, \Sigma) \quad (10)$$

in probability and in distribution, respectively. Thus $\widehat{X} \in \text{Sym}^+(p)$ is a consistent estimator of the log-Euclidean population mean (9) (Σ can be estimated consistently by the sample covariance matrix of the vectors $\text{vecd}(\log X_1), \dots, \text{vecd}(\log X_n)$ if finite fourth order moments are also assumed).

More interestingly, the log-Euclidean average \widehat{X} is asymptotically lognormal in the sense that, for large n , $\log \widehat{X}$ is approximately normally distributed in $\text{Sym}(p)$ with mean $\log M$ and covariance Σ/n . This leads to the first definition of the lognormal distribution.

Definition 1. We say that $X \in \text{Sym}^+(p)$ has a *PD-matrix-variate lognormal distribution of Type I* with parameters $M \in \text{Sym}^+(p)$ and $\Sigma \in \text{Sym}^+(q)$, denoted $X \sim LN_I(M, \Sigma)$, if $\log X \sim N(\log M, \Sigma)$.

Definition 1 has the interpretation that the parameter M is precisely the log-Euclidean mean (9) of the lognormal variable X . Based on (10), we can write the approximate distribution of the log-Euclidean average for large n as

$$\widehat{X} \sim LN_I(M, \Sigma/n). \quad (11)$$

To generate random lognormal PD matrices $X \sim LN_I(M, \Sigma)$, the recipe is to generate random normal symmetric matrices $Y \sim N(\log M, \Sigma)$ and then apply the matrix exponential $X = \exp(Y)$. An illustration for $p = 2$ is shown in Figure 2. The representation of the random matrices and their log-Euclidean mean and average as ellipses via (2) nicely visualizes their positive definiteness and variability. However, it is less useful for representing CR's, as explained below.

The following result helps establish the density of the PD-matrix-variate lognormal distribution of Type I. The proof is given in Appendix B.

Theorem 1. Let $\lambda_1 \geq \dots \geq \lambda_p > 0$ denote the eigenvalues of $X \in \text{Sym}^+(p)$. The Jacobian of the transformation $Y = \log X \in \text{Sym}(p)$ is given by

$$J(X) = \mathcal{J}(Y \rightarrow X) = \frac{1}{\lambda_1 \dots \lambda_p} \prod_{i < j} g(\lambda_i, \lambda_j), \quad (12)$$




Figure 2: (a) 30 samples (gray) generated from a lognormal distribution $X \sim LN(M, \Sigma)$ with parameters $M = \text{Diag}(4, 1)$ and $\Sigma = \text{Diag}(0.5, 0.5, 0)$ represented by ellipses via (2). Notice the variation around the mean ellipse M (red) with semiaxes $\sqrt{4} = 2$ and $\sqrt{1} = 1$. Superimposed are the log-Euclidean average (inner blue) and two extreme points of the 95% CR (outer blue). (b) Same as in (a) displayed in the coordinates of Figure 1, where each ellipse in (a) corresponds to a point of the same color inside the cone. The 95% CR is represented by the cyan mesh.

where

$$g(\lambda_i, \lambda_j) = \begin{cases} (\log \lambda_i - \log \lambda_j) / (\lambda_i - \lambda_j), & \lambda_i > \lambda_j \\ 1/\lambda_i, & \lambda_i = \lambda_j. \end{cases} \quad (13)$$

Making the change of variable $Y = \log X$ in Definition 1 and applying Theorem 1 gives that the density of $X \sim LN_I(M, \Sigma)$ is

$$f(X; M, \Sigma) = \frac{J(X)}{(2\pi)^{q/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2} \text{vecd}(\log X - \log M)' \Sigma^{-1} \text{vecd}(\log X - \log M)\right) \quad (14)$$

with respect to Lebesgue measure on $\text{Sym}^+(p)$, where $J(X)$ is the Jacobian (12). An interesting special case occurs when $M = aI_p$, $a > 0$, and Σ is an orthogonally invariant covariance matrix with respect to the operator $\text{vecd}(\cdot)$ (Schwartzman et al., 2008), e.g. $\Sigma = I_q$. In that case the density (14) becomes a function of the eigenvalues of X only and is orthogonally invariant in the sense that $f(X) = f(QXQ')$ for any $p \times p$ orthogonal matrix Q . We omit the explicit expression of the density in that case for brevity.

As an example of what the lognormal density looks like, Figure 3 shows the construction of a lognormal density of Type I in $\text{Sym}^+(2)$ with parameters

$$M = \exp \begin{pmatrix} 0.7 & 0.2 \\ 0.2 & 0 \end{pmatrix} = \begin{pmatrix} 2.05 & 0.29 \\ 0.29 & 1.03 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} 0.25 & 0.05 & 0 \\ 0.05 & 0.5 & 0 \\ 0 & 0 & 0.5 \end{pmatrix}. \quad (15)$$

Panel (a) shows two views of the density of the normal vector $y = \text{vecd}(Y) \sim N(\text{vecd}(\log M), \Sigma)$. Applying the transformation $X = \exp(Y)$ yields the density of the vector $x = \text{vecd}(X)$, shown in two views in Panel (b). The gridlines in Panel (a) have been mapped directly to Panel (b) in order to illustrate the deformative effect of the matrix exponential. Notice the shape of the $\text{Sym}^+(2)$ cone on the right column.




Figure 3: (a) Density of $y = \text{vecd}(Y) \sim N(\text{vecd}(\log M), \Sigma)$ with parameters (15) displayed in two planes: $y_3 = 0$ (left column) and $y_1 = y_2$ (right column). The density contours increase in height for darker colors. (b) Density of $x = \text{vecd}(X)$ with $X = \exp(Y)$ displayed in two planes: $x_3 = 0$ (left column) and $x_1 = x_2$ (right column).

Finally, another nice property of the lognormal distribution of Type I is that, if X_1, \dots, X_n are i.i.d. $LN(M, \Sigma)$, then the MLEs of M and Σ are given, respectively, by the log-Euclidean average and sample covariance matrix

$$\hat{M} = \hat{X} = \exp\left(\frac{1}{n} \sum_{i=1}^n \log X_i\right), \quad \hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \text{vecd}(\log X_i) \text{vecd}(\log X_i)'. \quad (16)$$

2.5 A confidence region for the log-Euclidean mean

Based on (11), a CR for the log-Euclidean mean of asymptotic level α is given by

$$\text{CR}_{n,\alpha}(\hat{X}) = \left\{ M \in \text{Sym}^+(p) : n \text{vecd}(\log \hat{X} - \log M)' \hat{\Sigma}^{-1} \text{vecd}(\log \hat{X} - \log M) \leq \chi_{q,1-\alpha}^2 \right\}, \quad (16)$$

where $\chi_{q,1-\alpha}^2$ is the $1-\alpha$ quantile of the χ^2 distribution with q degrees of freedom and $\hat{\Sigma}$ is a consistent estimator of Σ such as the sample covariance matrix of $\text{vecd}(\log X_1), \dots, \text{vecd}(\log X_n)$.

An example of such a CR is shown in Figure 2b. Note that the 3-dimensional ellipsoid determined by (16) in the log domain gets distorted after inverting the log into the original domain, so that the log-Euclidean average does not correspond to the Euclidean center of the CR but rather to its log-Euclidean center. The points on the CR mesh were obtained by solving for M in (16) using the principal components of $\hat{\Sigma}$ as basis. Particularly, marked in blue are the two extreme points along the first principal component of the CR (computed in the log

domain), computed by inverting (16) according to

$$\exp \left[\log \widehat{X} \pm \text{vecd}^{-1} \left(\lambda_1 V_1 \chi_{q,1-\alpha}^2 / n \right) \right],$$

where λ_1 , V_1 are the first eigenvalue and first eigenvector of the covariance matrix $\widehat{\Sigma}$. These two points are also represented as ellipses in Figure 2a. It is worth emphasizing that the shaded region in Figure 2a is not a CR because it only captures the variability along the first principal component. While the true mean is not contained in the shaded region in Figure 2a, it is indeed contained in the CR in Figure 2b. Other visualization schemes may also be devised, for instance, by choosing points on the boundary of the CR that are multiples of the log-Euclidean average or whose corresponding ellipses have the largest and smallest plotted area.

3 The canonical geometric average and the PD-matrix lognormal distribution of Type II

A second type of geometric average, called here canonical geometric average, is obtained as the Fréchet average according to the so-called canonical or affine-invariant metric defined on $\text{Sym}^+(p)$ (Arsigny et al., 2006; Lenglet et al., 2006; Fletcher and Joshi, 2007). The canonical geometric average also leads asymptotically to a lognormal distribution, but of a different kind.

3.1 The canonical geometric mean in $\text{Sym}^+(p)$

In general, the Fréchet mean of a probability measure Q on a complete metric space with distance ρ is the minimizer of the function $F(P) = \int \rho^2(P, X) Q(dX)$ (Fréchet, 1948; Pennec, 1999; Bhattacharya and Patrangenaru, 2003, 2005). In particular, for the empirical measure $\hat{Q}_n = (1/n) \sum_{i=1}^n \delta(X_i)$ of a set of i.i.d. data points X_1, \dots, X_n , the minimizer of the function $\hat{F}_n(P) = (1/n) \sum_{i=1}^n \rho^2(P, X_i)$ is called the sample Fréchet mean or Fréchet average.

Seen in this context, the Euclidean average (3) and corresponding Euclidean mean $E(X) = \int X Q(dX)$ are respectively the unique Fréchet average and mean according to the Euclidean metric defined on $\text{Sym}^+(p)$ and given by the Frobenius distance defined via

$$d_E^2(P, X) = |X - P|^2 = \text{tr}[(X - P)^2]. \quad (17)$$

Similarly, the log-Euclidean average (8) and log-Euclidean mean (9) are respectively the unique Fréchet average and mean according to a log-Euclidean distance defined on $\text{Sym}^+(p)$ defined via (Arsigny et al., 2006)

$$d_{LE}^2(P, X) = |\log X - \log P|^2 = \text{tr}[(\log X - \log P)^2]. \quad (18)$$

If the measure Q is defined on a manifold embedded in Euclidean space, then the Fréchet mean associated with the restriction to the manifold of the Euclidean distance in the ambient space is called the extrinsic mean. Because of the inclusion $\text{Sym}^+(p) \subset \text{Sym}(p)$, the embedding of $\text{Sym}^+(p)$ in Euclidean space is trivial with no change in dimension. Moreover, the embedding provided by the operator $\text{vecd}(\cdot)$ is isometric (see (30) in Appendix A.1) and so the Euclidean mean can be thought of as the extrinsic mean according to the Euclidean distance in \mathbb{R}^q , $q = p(p+1)/2$.

If the measure Q is defined on a Riemannian manifold with geodesic distance ρ , then the Fréchet mean associated with this distance is called the intrinsic mean, not requiring an embedding. In the case of $\text{Sym}^+(p)$, the canonical geometric mean of a measure Q is defined as

the intrinsic mean of Q according to the canonical geodesic distance defined via

$$d_C^2(P, X) = |\log(M^{-1/2} X M^{-1/2})|^2 = \text{tr} \{ [\log(M^{-1/2} X M^{-1/2})]^2 \}, \quad (19)$$

where the norm is again the Frobenius norm. With this metric, $\text{Sym}^+(p)$ becomes a geodesically complete Riemannian manifold (see Appendix A.4). The canonical geometric mean is characterized by the following result.

Proposition 1. *Let Q be a probability measure on $\text{Sym}^+(p)$ such that $F_C(P) = \int d_C^2(P, X) Q(dX)$ is finite for $P \in \text{Sym}^+(p)$, where d_C is the canonical geodesic distance (19). Then the canonical geometric mean M is the unique solution to the equation*

$$\int \log(M^{-1/2} X M^{-1/2}) Q(dX) = 0. \quad (20)$$

This result is a direct consequence of Theorem 2.1 of Bhattacharya and Patrangenaru (2003) for complete simply connected Riemannian manifolds. The uniqueness is guaranteed by the non-positive curvature of $\text{Sym}^+(p)$ with the canonical metric (Appendix A.4). Equation (20) indicates the desired property that the residuals about the intrinsic mean, defined here as the image of the data points onto the tangent space at M via the Riemannian logarithmic map (given by (40) in Appendix A.4) average to zero.

The canonical geometric mean has the following equivariance properties. If X is a random PD matrix according to a probability measure Q with canonical geometric mean M , then:

1. The canonical geometric mean of AXA^{-1} is AMA^{-1} , for $A \in GL(p)$.
2. The canonical geometric mean of GXG' is GMG' , for $G \in GL(p)$.
3. The canonical geometric mean of X^{-1} is M^{-1} .

These properties parallel the invariance properties of the canonical distance (see (41) in Appendix A.4). Note the additional equivariance with respect to the group action of $GL(p)$ (see (37) in Appendix A.4) not enjoyed by the log-Euclidean mean (Section 2.3).

3.2 The canonical geometric average in $\text{Sym}^+(p)$

For a set of i.i.d. data points $X_1, \dots, X_n \in \text{Sym}^+(p)$, their canonical geometric average, denoted here \hat{X} , is defined as their sample intrinsic mean according to the geodesic distance (19). Applying Proposition 1 to the empirical measure $\hat{Q}_n = (1/n) \sum_{i=1}^n \delta(X_i)$ gives that the canonical geometric average is unique and satisfies

$$\frac{1}{n} \sum_{i=1}^n \log(\hat{X}^{-1/2} X_i \hat{X}^{-1/2}) = 0. \quad (21)$$

This condition has been also derived independently by Moakher (2005).

Unlike the Euclidean and log-Euclidean averages, the canonical geometric average in general has no closed analytical form, except for a few cases detailed in Moakher (2005). The following iterative algorithm is commonly used for numerically computing the canonical geometric average (Pennec, 1999; Fletcher and Joshi, 2007).

Algorithm 1.

1. For $k = 0$, set an initial value $\mathring{X}^{(0)}$ such as the Euclidean average \bar{X} or the log-Euclidean average \widehat{X} . Fix $\epsilon > 0$.
2. Evaluate

$$\bar{Y}^{(k)} = \frac{1}{n} \sum_{i=1}^n \log \left[\left(\mathring{X}^{(k)} \right)^{-1/2} X_i \left(\mathring{X}^{(k)} \right)^{-1/2} \right], \quad \mathring{X}^{(k+1)} = \left(\mathring{X}^{(k)} \right)^{1/2} \exp \left(\bar{Y}^{(k)} \right) \left(\mathring{X}^{(k)} \right)^{1/2}.$$

3. If $|\bar{Y}^{(k)}| < \epsilon$, stop. Otherwise, increase k by 1 and go back to Step 2.

In Algorithm 1, each iteration in Step 2 maps the data onto the tangent space at $\mathring{X}^{(k)}$ via the Riemannian log map, finds the Euclidean average $\bar{Y}^{(k)}$ there, and then maps it back onto the manifold via the Riemannian exponential map to obtain $\mathring{X}^{(k+1)}$. The stopping condition in Step 3 reflects the property (21) that the average of the residuals should converge to zero.

Because the canonical geometric average is a special case of the canonical geometric mean with the empirical measure \hat{Q}_n , it has the same equivariance properties. Namely, if X_1, \dots, X_n are PD matrices with canonical geometric average \mathring{X} , then:

1. The canonical geometric average of $AX_1A^{-1}, \dots, AX_nA^{-1}$ is $A\mathring{X}A^{-1}$, for $A \in GL(p)$.
2. The canonical geometric average of GX_1G', \dots, GX_nG' is $G\mathring{X}G'$, for $G \in GL(p)$.
3. The canonical geometric average of $X_1^{-1}, \dots, X_n^{-1}$ is \mathring{X}^{-1} .

The consistency of the sample canonical geometric mean is given by the following result.

Proposition 2. *Let Q be a probability measure on $\text{Sym}^+(p)$ such that $F_C(P) = \int d_C^2(P, X) Q(dX)$ is finite for $P \in \text{Sym}^+(p)$, where d_C is the canonical geodesic distance (19). Then the canonical geometric average \mathring{X} is a strongly consistent estimator of the canonical geometric mean M .*

This result is a direct consequence of Theorem 2.3 of Bhattacharya and Patrangenaru (2003) for complete metric spaces and is understood in the sense that, given any $\epsilon > 0$, the sample intrinsic mean will be within a distance less than ϵ from the true intrinsic mean for all sufficiently large sample sizes with probability 1.

3.3 The canonical geometric average and the PD-matrix lognormal distribution of Type II

Like in the case of the log-Euclidean average, the asymptotic distribution of the canonical geometric average also leads to a matrix lognormal distribution on $\text{Sym}^+(p)$. To obtain this asymptotic distribution, we rely on the CLT for Riemannian manifolds given by Theorem 2.2 of Bhattacharya and Patrangenaru (2005), further discussed in Bhattacharya and Bhattacharya (2008).

Theorem 2. *Let $X_1, \dots, X_n \in \text{Sym}^+(p)$ be i.i.d. PD matrices according to a probability measure Q that is absolutely continuous with respect to the Euclidean volume measure in $\text{Sym}(p)$. Let M and \mathring{X} be the canonical geometric mean and canonical geometric average, respectively. Further, define the i.i.d. random symmetric matrices*

$$Y_i = \log \left(M^{-1/2} X_i M^{-1/2} \right), \quad i = 1, \dots, n \quad (22)$$

and assume their common covariance matrix $\Sigma = \text{Cov}(\text{vecd}(Y_1))$ is finite. Then,

$$\sqrt{n} \text{vecd} \left[\log \left(M^{-1/2} \mathring{X} M^{-1/2} \right) \right] \Rightarrow N(0, K^{-1} \Sigma K^{-1}) \quad (23)$$

in distribution, with

$$K = I + D' E \left[\frac{1}{12} (Y_1 \ominus Y_1)^2 + \frac{1}{720} (Y_1 \ominus Y_1)^4 + \dots \right] D, \quad (24)$$

where D is the $p^2 \times q$ duplication matrix defined by (31) in Appendix A.1, and the symbol “ \ominus ” denotes the operation $Y \ominus Y = Y \otimes I - I \otimes Y$.

The operator “ \ominus ” may be called “Kronecker difference” in resemblance with the Kronecker sum operator $Y \oplus Y = Y \otimes I + I \otimes Y$. The matrix K reflects the effect of curvature on the asymptotic covariance. The proof of Theorem 2 is given in Appendix B.

As with the arithmetic and log-Euclidean averages, the limiting distribution of the canonical geometric average is defined on the tangent space at the true value of the canonical geometric mean. Theorem 2 states that the canonical geometric average \mathring{X} is asymptotically lognormal in the sense that, for large n , $\log \left(M^{-1/2} \mathring{X} M^{-1/2} \right)$ is approximately normally distributed in $\text{Sym}(p)$ with mean zero and covariance $K^{-1} \Sigma K^{-1}/n$. This leads to the second definition of the lognormal distribution.

Definition 2. We say that $X \in \text{Sym}^+(p)$ has a *PD-matrix-variate lognormal distribution of Type II* with parameters $M \in \text{Sym}^+(p)$ and $\Sigma \in \text{Sym}^+(q)$, denoted $X \sim LN_{II}(M, \Sigma)$, if $\log \left(M^{-1/2} X M^{-1/2} \right) \sim N(0, \Sigma)$.

Definition 2 has the interpretation that the parameter M is precisely the canonical geometric mean of the lognormal variable X because it satisfies (20). Based on (23), Definition 2 allows us to write the approximate distribution of the log-Euclidean average for large n as

$$\mathring{X} \sim LN_{II}(M, K^{-1} \Sigma K^{-1}/n). \quad (25)$$

Note that the lognormal distribution given by Definition 2 is not the same as the lognormal distribution given by Definition 1. This is because, in general, $\log \left(M^{-1/2} X M^{-1/2} \right) \neq \log X - \log M$. The only case where equality holds for all $X \in \text{Sym}^+(p)$ is when M is a multiple of the identity matrix, i.e. $M = aI$ for some $a > 0$. In this case, the lognormal distributions of Type I and II coincide.

From Definition 2, the density of $X \sim LN_{II}(M, \Sigma)$ is obtained by the change of variable $Y = \log(M^{-1/2} X M^{-1/2})$. Applying Theorem 1 to the random variable $\tilde{X} = M^{-1/2} X M^{-1/2}$, in addition to the Jacobian $J(\tilde{X} \rightarrow X) = 1/|M|$, gives that X has density

$$f(X; M, \Sigma) = \frac{J(\tilde{X})}{(2\pi)^{q/2} |\Sigma|^{1/2} |M|} \exp \left(-\frac{1}{2} \text{vecd}(\log \tilde{X})' \Sigma^{-1} \text{vecd}(\log \tilde{X}) \right) \quad (26)$$

with respect to Lebesgue measure on $\text{Sym}^+(p)$, where $J(\tilde{X})$ is given by expression (12) but replacing instead the eigenvalues $\tilde{\lambda}_1 \geq \dots \geq \tilde{\lambda}_p > 0$ of $\tilde{X} = M^{-1/2} X M^{-1/2}$, also known as the generalized eigenvalues of the pair (M, X) .

For illustration, Figure 4 shows the construction of a lognormal density of Type II in $\text{Sym}^+(2)$ with the same parameters as (15). Panel (a) shows two views of the density of the normal vector $y = \text{vecd}(Y) \sim N(0, \Sigma)$. Applying the transformation $X = \exp(M^{1/2} Y M^{1/2})$ yields the density of the vector $x = \text{vecd}(X)$, shown in two views in Panel (b). The gridlines in Panel (a) have been mapped directly to Panel (b) in order to illustrate the deformative effect of the exponential transformation, which in contrast to Figure 3, is affected by M . Notice the shape of the $\text{Sym}^+(2)$ cone on the right column and the additional deformation in comparison with Figure 3.




Figure 4: (a) Density of $y = \text{vecd}(Y) \sim N(0, \Sigma)$ with parameters (15) displayed in two planes: $y_3 = 0$ (left column) and $y_1 = y_2$ (right column). The density contours increase in height for darker colors. (b) Density of $x = \text{vecd}(X)$ with $X = \exp(M^{1/2}Y M^{1/2})$ displayed in two planes: $x_3 = 0$ (left column) and $x_1 = x_2$ (right column).

3.4 A confidence region for the canonical geometric mean

Based on (23), a confidence region for the canonical geometric mean of asymptotic level α is given by

$$\text{CR}_{n,\alpha}(\hat{X}) = \{M \in \text{Sym}^+(p) : n \text{vecd}(\hat{Y})' \hat{K} \hat{\Sigma}^{-1} \hat{K} \text{vecd}(\hat{Y}) \leq \chi_{q,1-\alpha}^2\} \quad (27)$$

where

$$\hat{Y} = \log(M^{-1/2} \hat{X} M^{-1/2})$$

The matrices $\hat{\Sigma}$ and \hat{K} in (27) are consistent estimators of Σ and K , respectively, and may be obtained as follows. Letting

$$\hat{Y}_i = \log(\hat{X}^{-1/2} X_i \hat{X}^{-1/2}), \quad i = 1, \dots, n$$

in accordance to (22) where \hat{X} estimates M , we can take

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \text{vecd}(\hat{Y}_i) \text{vecd}(\hat{Y}_i)'$$

Note that there is no need to subtract the average of the $\text{vecd}(\hat{Y}_i)$'s in the computation of the sample covariance because it is zero by (21). Similarly, from (24), we can take

$$\hat{K} = I + D' \frac{1}{n} \sum_{i=1}^n \left[\frac{1}{12} (\hat{Y}_i \ominus \hat{Y}_i)^2 + \frac{1}{720} (\hat{Y}_i \ominus \hat{Y}_i)^4 + \dots \right] D.$$

It is worth noting that in (27), \mathring{Y} depends on M in a nonlinear way. In the data analysis Section 4 below we will compute the extreme points along the first principal component of the CR (in the log domain). These are defined as the points M satisfying the equation

$$M^{-1/2} \mathring{X} M^{-1/2} = B_{\pm} \quad (28)$$

where

$$B_{\pm} = \exp \left[\pm \text{vecd}^{-1} \left(\lambda_1 V_1 \chi_{q,1-\alpha}^2 / n \right) \right]$$

and λ_1 , V_1 are the first eigenvalue and first eigenvector of the matrix $\hat{K}^{-1} \hat{\Sigma} \hat{K}^{-1}$. Equation (28) is a special case of the continuous-time algebraic Ricatti equation. Interestingly, its solution is also a canonical geometric average in the following sense. By Lawson and Lim (2001), the unique solution to the equation $XA^{-1}X = B$ for $A, B \in \text{Sym}^+(p)$ is the canonical geometric average of A and B defined via (21) which, for $n = 2$, has the closed form solution $A \# B = A(A^{-1}B)^{1/2}$ (Moakher, 2005). Therefore the solution to (28) is

$$M = \left[(\mathring{X}B)^{-1/2} \mathring{X} \right]^2.$$

4 A DTI data example

4.1 Data description

Diffusion Tensor Imaging (DTI) is a particular modality of magnetic resonance imaging often used to visualize the brain's white matter (Basser and Pierpaoli, 1996; LeBihan et al., 2001). DTI images are 3D rectangular arrays that contain at every voxel a 3×3 symmetric PD matrix, also called diffusion tensor (DT). A useful scalar summary of the DT is fractional anisotropy (FA), a function of the DT's eigenvalues. FA values near 0 represent nearly isotropic diffusion, found in the brain's water ventricles, while values near 1 represent highly anisotropic diffusion, found in tightly packed white matter neural fibers. Another useful summary is the principal diffusion direction (PDD), the first eigenvector of the DT. In highly anisotropic voxels it indicates the direction of the underlying neural fibers.

In this paper, we analyze a dataset consisting of 34 DTI images corresponding to 34 10-year-old children (12 boys and 22 girls). This dataset was also analyzed in Schwartzman et al. (2010) and is part of a larger observational study of brain anatomy in children (Dougherty et al., 2007). The 34 DTI images were spatially normalized to a common template (Dougherty et al., 2005), resulting in each image being a voxel array of size $81 \times 106 \times 76$ in a rectangular grid with $2 \times 2 \times 2$ mm regular spacings. The coordinate axes are defined so that the x , y and z axis point respectively to the right, front and top of the brain (Figure 5). Like in Schwartzman et al. (2010), analysis was restricted to voxels strictly inside the brain.

4.2 Data analysis

After spatial normalization, the collection of 34 DT's across subjects at each voxel can be modeled as an i.i.d. sample of PD matrices whose population mean and covariance depends on the voxel location. As such, estimation of the population mean DT image is achieved by computing PD averages voxelwise (Schwartzman et al., 2010).

Figure 5 shows the three types of voxelwise averages discussed in this paper at a transverse slice 36 mm above the anterior commissure, an anatomical landmark commonly used for spatial normalization (Talairach convention). This is the same slice that is shown in Schwartzman




Figure 5: Voxelwise DT averages. Image intensity is proportional to FA in the range $[0, 0.6]$. Colors indicate coordinate directions of the PDD: right-left (red), anterior-posterior (green) and superior-inferior (blue). Mixed colors represent directions that are oblique to the coordinate axes.

et al. (2010). Because DTs are 6-dimensional objects, Figure 5 shows a meaningful reduction to 3 dimensions. One dimension is captured by FA, encoded in the image intensity, while the other two are captured by the PDD, encoded in color.

At first, all three averages appear to be very similar. However, Figure 6 shows that the Euclidean and log-Euclidean averages disagree in some regions, particularly in the corpus callosum (red X-like structure in Figure 5). The differences between the log-Euclidean and canonical geometric averages are about an order of magnitude smaller.

In the current debate about average types, the DTI literature has not yet considered whether the observed differences between them may be due to sampling variability. Constructing a formal two-sample test for a pair of means is difficult because they are defined by different geometries and because of the dependence between them, as they are based on the same data. More informally, to make use of the CRs developed above, we may check whether the value of a particular type of average is inside the CR of another type of average. This approach is conservative in detecting differences because the variability in any type of average is smaller than the variability in the difference between two types of averages.

Following this approach, Figure 7a evaluates the log-Euclidean average with respect to the Euclidean average. Shown is a map of the smallest α such that the log-Euclidean average \widehat{X} is inside the Euclidean CR defined by (7). This may be loosely interpreted as a p-value map for the observed log-Euclidean average under the assumption that the Euclidean average is the true mean. Interestingly, only 28 out of the 3419 voxels in this slice have a p-value smaller than 0.05, and only 4 of them have a p-value smaller than the 2.3×10^{-4} threshold corresponding to an FDR of 0.2 using the Benjamini-Hochberg algorithm (Benjamini and Hochberg, 1995).

To double-check, Figure 7b does the reverse and evaluates the Euclidean average with respect to the log-Euclidean average by showing a map of the smallest α such that the Euclidean average \bar{X} is inside the log-Euclidean CR defined by (16). This may be interpreted loosely as a p-value map for the observed Euclidean average under the assumption that the log-Euclidean average is the true mean. In this comparison, only 16 out of the 3419 voxels in this slice have a p-value smaller than 0.05, and none of them survive the 0.2 FDR cutoff. In summary, most of the differences between the Euclidean and log-Euclidean averages are indeed within the range of random variation.




Figure 6: (a) FA differences between: log-Euclidean and Euclidean average (left); log-Euclidean and canonical (right). Image intensity scale is the same for both panels. (b) Angle between: PDD of log-Euclidean and Euclidean average (left); PDD of log-Euclidean and canonical (right). Image intensity scale in degrees is the same in both panels.

For completeness, Figure 7c shows the corresponding p-value map comparing the canonical geometric average to the log-Euclidean CR. Here all p-values are very close to 1. The results are similar when comparing the log-Euclidean average to the canonical geometric CR (map not shown). In other words, the log-Euclidean and canonical geometric averages are statistically indistinguishable.

To better visualize the few significant differences between the three types of averages, Figure 8 shows their ellipsoidal rendering via (2) for the voxel with the smallest p-value in Figure 7a. Superimposed are similar ellipsoidal renderings of the extreme points along the first principal component (in the log domain) of the corresponding 95% CR, as was done for the 2D example in Figure 2. The observed averages do exhibit visible differences in FA and orientation, but so do the extreme points of the CRs. Note that these 6-dimensional extreme points do not show the full extent of the 21-dimensional CRs but only their main mode of variation.

5 Discussion

5.1 Metrics and averages in $\text{Sym}^+(p)$

In this paper we have considered the three most common types of averages and their corresponding population means defined on $\text{Sym}^+(p)$: arithmetic or Euclidean, log-Euclidean and canonical geometric. We have seen that the CLT applied to the latter two leads to the definition




Figure 7: Log_{10} p-values of voxelwise DT averages: (a) Log-Euclidean average with respect to Euclidean CR; (b) Euclidean average with respect to log-Euclidean CR; (c) Canonical geometric average with respect to log-Euclidean CR. Image intensity scale is the same in all panels.




Figure 8: The voxel with the largest differences between the average types (brightest voxel in Figure 7a): ellipsoidal rendering of the DT averages (red) and two extreme points of the corresponding 95% CR (blue). Colored axes correspond to the coordinate directions described in Figure 5.

of a lognormal distribution on $\text{Sym}^+(p)$, analogous to the lognormal distribution for positive scalars.

These three types of averages correspond to three different metrics defined on $\text{Sym}^+(p)$. The preference for a particular metric and corresponding average type may be often based on qualitative interpretation of the results in DTI analysis (Fletcher and Joshi, 2007; Whitcher et al., 2007; Dryden et al., 2009; Pasternak et al., 2010). Often for data on manifolds, such as data on the sphere, Riemannian metrics are appealing because they lead to intrinsic means instead of extrinsic means. Computational considerations are also important (Groisser, 2004; Bhattacharya et al., 2012). However, in the case of $\text{Sym}^+(p)$ with the Euclidean metric, the intrinsic and extrinsic means coincide and are equal to the simple Euclidean average, which has a well defined asymptotic normal distribution despite $\text{Sym}^+(p)$ not being a vector space. From this point of view, going beyond Euclidean metrics for $\text{Sym}^+(p)$ may seem unnecessary.

In this paper, we have considered the random variability of each average. According to the data analysis of Section 4, the log-Euclidean and canonical geometric averages are virtually indistinguishable, while the Euclidean and log-Euclidean averages are mostly within the expected random variation for the given sample size. When comparing their computational complexity, our Matlab implementation computed the Euclidean, log-Euclidean and canonical geometric averages for all 3419 voxels in the slice shown in 0.72 s, 3.48 s and 16.48 s, respectively. These results make the log-Euclidean average convincingly preferable to the canonical geometric average. The choice between the Euclidean and log-Euclidean averages may be based on their desired equivariance properties, but perhaps is only important for larger sample sizes. The statistical relevance of other types of averages suggested by Dryden et al. (2009) could be analyzed using large sample techniques like the ones used in this paper.

5.2 Parametric models for random PD matrices

The two PD-matrix-variate lognormal families of distributions defined in this paper were obtained by applying the CLT to the log-Euclidean and canonical geometric averages of PD matrices. We derived their explicit densities and provided CRs for the corresponding geometric means.

More generally, the PD-matrix-variate lognormal families could be used as parametric models for PD matrix data, either for frequentist data analysis or as flexible priors for Bayesian analysis. The lognormal distributions provide a more flexible parametric model than the Wishart distribution because it has $q = p(p + 1)/2$ parameters for the mean and $q(q + 1)/2$ separate additional parameters for the covariance, while the Wishart distribution has q parameters for the mean and one parameter for the number of degrees of freedom, which affects both the covariance and the mean. The matrix lognormal distribution, on the other hand, is heavier-tailed than the Wishart, just like the scalar lognormal is heavier-tailed than the χ^2 distribution.

Going beyond PD matrices, the concept of a matrix-variate lognormal distribution could be extended to other Riemannian manifolds where the Riemannian logarithmic map is the matrix logarithm. One example is the orthogonal group $O(p)$, whose tangent space is the set of antisymmetric or skew-symmetric matrices. Extrinsic and intrinsic averages on that set have been studied by Moakher (2002). Applying the CLT to the log-Euclidean average of orthogonal matrices should lead naturally to a lognormal distribution on that set. The details are left for future work.

Acknowledgments

Special thanks to Emmanuel Ben-David (Columbia U.), Victor Patrangenaru (Florida State U.) and David Groisser (U. of Florida, Gainesville) for helpful discussions, as well as Bob Dougherty (Stanford U.) for providing the DTI data. This work was partially funded by NIH grant R21EB012177.

A Auxiliary definitions and results

A.1 The $\text{vecd}(\cdot)$ operator

For a matrix $Y \in \text{Sym}(p)$, we define the mapping of $\text{Sym}(p)$ to \mathbb{R}^q with $q = p(p+1)/2$ given by

$$\text{vecd}(Y) = (\text{diag}(Y)', \sqrt{2} \text{offdiag}(Y)')', \quad (29)$$

where $\text{diag}(Y)$ is a $p \times 1$ column vector containing the diagonal entries of Y and $\text{offdiag}(Y)$ is a $(q-p) \times 1$ column vector containing the off-diagonal entries of Y copied from below the diagonal columnwise (or above the diagonal rowwise) (Schwartzman et al., 2010). For example, in the case $p = 3$, $q = 6$,

$$\text{vecd}(Y) = (Y_{11}, Y_{22}, Y_{33}, \sqrt{2} Y_{12}, \sqrt{2} Y_{13}, \sqrt{2} Y_{23})'.$$

In contrast to the usual columnwise vectorization operator $\text{vec}(\cdot)$ (Gupta and Nagar, 2000; Chiu et al., 1996), $\text{vecd}(\cdot)$ provides a more natural representation for data in the form of symmetric matrices and it has the convenient property that the Frobenius norm of Y is the same as the Euclidean norm of $\text{vecd}(Y)$, i.e.

$$\text{tr}(Y^2) = \text{vecd}(Y)' \text{vecd}(Y). \quad (30)$$

The operator $\text{vecd}(\cdot)$ provides an embedding of both $\text{Sym}(p)$ and $\text{Sym}^+(p)$ in \mathbb{R}^q . This is the same embedding illustrated for $p = 2$ in Figure 1. The operator $\text{vecd}(\cdot)$ is invertible and its inverse shall be denoted by $\text{vecd}^{-1}(\cdot)$.

For any p , a duplication matrix D of size $p^2 \times q$ exists such that the operator $\text{vecd}(\cdot)$ is related to the usual columnwise vectorization operator $\text{vec}(\cdot)$ by

$$\text{vec}(Y) = D \text{vecd}(Y), \quad \text{vecd}(Y) = D' \text{vec}(Y). \quad (31)$$

For example, in the case $p = 3$, $q = 6$,

$$D = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/\sqrt{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/\sqrt{2} & 0 \\ 0 & 0 & 0 & 1/\sqrt{2} & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/\sqrt{2} \\ 0 & 0 & 0 & 0 & 1/\sqrt{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/\sqrt{2} \\ 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}.$$

The duplication matrix D has rank q and it satisfies the following properties. For any $Y, A \in \text{Sym}(p)$,

$$\begin{aligned} D'D &= I \\ DD'\text{vec}(Y) &= \text{vec}(Y) \\ \text{vecd}(AYA') &= D'(A \otimes A)D\text{vecd}(Y), \end{aligned}$$

where the third property is based on the general property that $\text{vec}(AYB') = (B \otimes A)\text{vec}(Y)$ (Schott, 2005).

A.2 Matrix exponential and logarithm

The exponential of a square matrix Y is defined as $X = \exp(Y) = e^Y = \sum_{k=0}^{\infty} Y^k/k!$. The inverse operation $Y = \log(X)$ is called the matrix logarithm of X . In particular, if $Y \in \text{Sym}(p)$ and $Y = VLV'$ is an eigen-decomposition with $V \in O(p)$ and $L \in \text{Diag}(p)$, then it follows from the series expansion that

$$X = \exp(Y) = V \exp(L) V', \quad (32)$$

where $\exp(L)$ is a diagonal matrix containing the exponential of each of the diagonal elements of L . Since the diagonal elements of L are real, the diagonal elements of $\exp(L)$ are positive, implying that X is PD. Given $X \in \text{Sym}^+(p)$, (32) can be inverted and the matrix log of X is uniquely given by

$$Y = \log(X) = V \log(\Lambda) V',$$

where $\log(\Lambda)$ is a diagonal matrix containing the logs of each of the diagonal elements of Λ .

Both $\exp(\cdot)$ and $\log(\cdot)$ are analytical functions and enjoy the useful properties with respect to similarity transformations and matrix inversion

$$\begin{aligned} \exp(AYA^{-1}) &= A \exp(Y) A^{-1}, & (e^Y)^{-1} &= e^{-Y} \\ \log(AXA^{-1}) &= A \log(X) A^{-1}, & \log(X^{-1}) &= -\log(X) \end{aligned} \quad (33)$$

for $A \in GL(p)$, the set of $p \times p$ real invertible matrices. In general, $e^{Y+Z} \neq e^Y e^Z$ and $\log(XW) \neq \log(X) + \log(W)$ unless Y and Z (or X and W) commute, i.e. they have the same eigenvectors. Instead, there is the Baker-Campbell-Hausdorff (BCH) formula

$$\log(e^Y e^Z) = Y + Z + \frac{1}{2}[Y, Z] + \frac{1}{12}[Y, [Y, Z]] + \frac{1}{12}[Z, [Z, Y]] + \frac{1}{24}[[Y, [Y, Z]], Z] + \dots$$

where $[Y, Z] = YZ - ZY$ is the matrix commutator. Notice that indeed $\log(e^Y e^Z) = Y + Z$ if Y and Z commute. In this paper we use an explicit polynomial expression without the commutators, also known as Goldberg's expansion (Goldberg, 1956). Using the method of

Newman and Thompson (1987), the first several terms of this series are found to be

$$\begin{aligned}
\log(e^Y e^Z) = & Y + Z + \frac{1}{2}(YZ - ZY) + \frac{1}{12}(YZ^2 + Y^2Z + ZY^2 + Z^2Y) - \frac{1}{6}(YZY + ZYZ) \\
& + \frac{1}{24}(Y^2Z^2 - Z^2Y^2) - \frac{1}{12}(YZYZ - ZYZY) - \frac{1}{720}(YZ^4 + Y^4Z + ZY^4 + Z^4Y) \\
& + \frac{1}{180}(Y^2Z^3 + Y^3Z^2 + Z^2Y^3 + Z^3Y^2 + YZY^3 + YZ^3Y + Y^3ZY + ZYZ^3 + ZY^3Z + Z^3YZ) \\
& - \frac{1}{120}(YZ^2Y^2 + Y^2ZY^2 + Y^2Z^2Y + ZY^2Z^2 + Z^2YZ^2 + Z^2Y^2Z) \\
& - \frac{1}{120}(YZYZ^2 + YZY^2Z + YZ^2YZ + Y^2ZYZ + ZYZY^2 + ZYZ^2Y + ZY^2ZY + Z^2YZY) \\
& + \frac{1}{30}(YZYZY + ZYZYZ) + \dots
\end{aligned} \tag{34}$$

A.3 $\text{Sym}^+(p)$ as a Riemannian manifold with the Euclidean and log-Euclidean metrics

In the view of Section 2.2, $\text{Sym}^+(p)$ is seen as a Riemannian submanifold of $\text{Sym}(p)$ endowed with the Euclidean or Frobenius metric inherited from $\text{Sym}(p)$. Given two tangent vectors $Y, Z \in \text{Sym}(p)$, their Euclidean or Frobenius inner product and associated squared norm are defined as

$$\langle Y, Z \rangle = \text{tr}(YZ) = \text{vecd}(Y)' \text{vecd}(Z), \quad |Y|^2 = \text{tr}(Y^2), \tag{35}$$

and the geodesic distance between two points $X_1, X_2 \in \text{Sym}^+(p)$ is the Euclidean distance defined via (17).

A Riemannian manifold is called geodesically complete if at any point on the manifold, the Riemannian exponential map, which maps tangent vectors to the manifold along geodesics, is defined for all tangent vectors. This property depends on the metric that is defined on the manifold. When endowed with the Euclidean metric, $\text{Sym}^+(p)$ is not geodesically complete because the Riemannian exponential map from the tangent space $\text{Sym}(p)$ to $\text{Sym}^+(p)$ at any point, which for the Euclidean metric is equal to a displacement by a tangent vector, is only defined for those tangent vectors whose displacement falls within the boundaries of $\text{Sym}^+(p)$. As a consequence, the approximate normal distribution of the Euclidean average (5) is defined on $\text{Sym}(p)$ but not on $\text{Sym}^+(p)$.

One way of achieving geodesic completeness is via the global matrix log transformation (1), which effectively removes the PD constraints by mapping each element of $\text{Sym}^+(p)$ to a unique element of $\text{Sym}(p)$ (Arsigny et al., 2006; Schwartzman, 2006). This is the view of Sections 2.3 and 2.4. According to this metric, the geodesic distance between two points $X_1, X_2 \in \text{Sym}^+(p)$ is the log-Euclidean distance defined via (18). Because of geodesic completeness, the approximate lognormal distribution of the log-Euclidean average (11) is well defined on $\text{Sym}^+(p)$.

As shown in Sections 2.3 and 2.4, the log transformation is simple enough that it permits obtaining results without the need of the Riemannian machinery.

A.4 $\text{Sym}^+(p)$ as a Riemannian manifold with the canonical metric

A more elaborate way of producing a geodesically complete manifold is via the so-called canonical or affine-invariant metric. For tangent vectors $Y, Z \in \text{Sym}(p)$ at $M \in \text{Sym}^+(p)$, their

canonical or affine invariant inner product and associated squared norm are defined as

$$\langle\langle Y, Z \rangle\rangle_M = \text{tr}(M^{-1}YM^{-1}Z), \quad \|Y\|_M^2 = \text{tr}[(M^{-1}Y)^2]. \quad (36)$$

Notice that at $M = I$, the $p \times p$ identity matrix, the canonical inner product (36) reduces to the Frobenius inner product (35). Various aspects of the geometry of $\text{Sym}^+(p)$ under the metric (36) have been described elsewhere (Kobayashi and Nomizu, 1996; Lang, 1999; Moakher, 2005; Smith, 2005; Schwartzman, 2006; Lenglet et al., 2006; Zhu et al., 2009), which we summarize here for completeness and easier reference.

Define the linear group action of $GL(p)$ on $\text{Sym}^+(p)$

$$\phi(G, X) = GXG', \quad G \in GL(p), \quad X \in \text{Sym}^+(p), \quad (37)$$

which maps the PD matrix X to the PD matrix GXG' . The metric (36) is called affine-invariant because it has the isometric property of being invariant under the group action (37), that is

$$\langle\langle Y, Z \rangle\rangle_M = \langle\langle GYG', GZG' \rangle\rangle_{GMG'} \quad (38)$$

for any $G \in GL(p)$. In particular, if $G = M^{-1/2}$, then (38) becomes

$$\langle\langle Y, Z \rangle\rangle_M = \langle M^{-1/2}YM^{-1/2}, M^{-1/2}ZM^{-1/2} \rangle_I = \text{tr}(M^{-1/2}YM^{-1/2} \cdot M^{-1/2}ZM^{-1/2}).$$

This says that the affine-invariant inner product can be computed by translating the tangent vectors to the tangent space at the identity and computing the standard Frobenius inner product there.

This device is useful for obtaining the Riemannian exponential and logarithmic maps corresponding to the affine-invariant metric. Since the Riemannian exponential map of a tangent vector $Y \in \text{Sym}(p)$ at the identity I is equal to the matrix exponential $\exp(Y) \in \text{Sym}^+(p)$, the Riemannian exponential map of a tangent vector Y at M onto $\text{Sym}^+(p)$ is obtained by mapping Y to the tangent space at the identity via $\phi(M^{-1/2}, \cdot)$, evaluating the Riemannian exponential map there, and mapping the result back again via $\phi(M^{1/2}, \cdot)$. The Riemannian exponential map is thus given by

$$X = \text{Exp}_M(Y) = M^{1/2} \exp(M^{-1/2}YM^{-1/2})M^{1/2} = M \exp(M^{-1}Y), \quad (39)$$

where the second expression is obtained via (33). By (1), this map is one-to-one. The inverse map is the Riemannian logarithmic map of $X \in \text{Sym}^+(p)$ to the tangent vector $Y \in \text{Sym}(p)$ in the tangent space at M , and is given by

$$Y = \text{Log}_M(X) = M^{1/2} \log(M^{-1/2}XM^{-1/2})M^{1/2} = M \log(M^{-1}X). \quad (40)$$

Because the Riemannian exponential map is bijective, there is a unique geodesic path joining any two points in $\text{Sym}^+(p)$. The geodesic distance between them may be defined uniquely as the length of that path, which by Gauss' lemma (DoCarmo, 1992, p. 69) is equal to the length of the projection via the Riemannian logarithmic map of one of the points onto the tangent space at the other. For any two points $M, X \in \text{Sym}^+(p)$, their geodesic distance according to the affine-invariant metric (36) is defined via (19).

The bijective property of the Riemannian exponential map makes $\text{Sym}^+(p)$ geodesically complete. Consequently, it is also complete as a metric space by the Hopf-Rinow theorem (DoCarmo, 1992, p. 147). The geodesic distance (19) is a proper metric and satisfies the properties of semi-positiveness, symmetry and the triangle inequality. In addition, it is invariant

under similarity transformations, under the group action (37) and under inversion (Forstner and Moonen, 1999):

$$\begin{aligned} d_C(M, X) &= d_C(AMA^{-1}, AXA^{-1}), & \forall A \in GL(p) \\ d_C(M, X) &= d_C(GMG', GYG'), & \forall G \in GL(p) \\ d_C(M, X) &= d_C(M^{-1}, X^{-1}). \end{aligned} \tag{41}$$

These properties make $\text{Sym}^+(p)$ a Riemannian homogeneous space and a symmetric space in the sense of Cartan (Kobayashi and Nomizu, 1996, Ch. X-XI).

The curvature of $\text{Sym}^+(p)$ under the affine-invariant inner product is non-positive and has been derived explicitly by Lenglet et al. (2006). Intuitively, non-positive curvature means that geodesics “diverge” away from each other. This property guarantees the existence and uniqueness of the canonical geometric mean (Proposition 1).

B Proofs

B.1 Proof of Theorem 1

Let $X = V\Lambda V'$ and $Y = VLV'$ be eigen-decompositions of X and Y so that $L = \log \Lambda$ with diagonal entries $l_i = \log \lambda_i$. Suppose the eigenvalues are all distinct. The Jacobian of the eigenvalue decomposition $X \rightarrow (V, \Lambda)$ can be written as $\mathcal{J}(X \rightarrow (V, \Lambda)) = g(V) \prod_{i < j} (\lambda_i - \lambda_j)$, where $g(V)$ is a function of the eigenvectors only (Mehta, 1991, p. 58). Similarly, $\mathcal{J}(Y \rightarrow (V, L)) = g(V) \prod_{i < j} (l_i - l_j)$. Finally, since $\partial l_i / \partial \lambda_i = 1/\lambda_i$, the transformation $L = \log \Lambda$ has Jacobian $\mathcal{J}(L \rightarrow \Lambda) = |\Lambda|^{-1}$. Putting the three Jacobians together,

$$\begin{aligned} \mathcal{J}(Y \rightarrow X) &= \mathcal{J}(Y \rightarrow (V, L)) \cdot \mathcal{J}((V, L) \rightarrow (V, \Lambda)) \cdot \mathcal{J}((V, \Lambda) \rightarrow X) \\ &= \left(g(V) \prod_{i < j} (l_i - l_j) \right) |\Lambda|^{-1} \left(g(V) \prod_{i < j} (\lambda_i - \lambda_j) \right)^{-1} = \frac{1}{\lambda_1 \dots \lambda_p} \prod_{i < j} \frac{\log \lambda_i - \log \lambda_j}{\lambda_i - \lambda_j}. \end{aligned}$$

The case $\lambda_i = \lambda_j$ is obtained by taking the limit of the fraction inside the product sign above as $\lambda_i \rightarrow \lambda_j$.

B.2 Proof of Theorem 2

B.2.1 Asymptotic normality

By Remark 2.2 of Bhattacharya and Patrangenaru (2005), Theorem 2.2 of Bhattacharya and Patrangenaru (2005) regarding the CLT for sample intrinsic means applies to $\text{Sym}^+(p)$ with the canonical inner product because it has no cut-locus and the supremum of its sectional curvatures is nonpositive. Note that the measure Q need not have compact support, as long as the second moment conditions in the statement of Theorem 2 are satisfied. Applying Theorem 2.2 of Bhattacharya and Patrangenaru (2005) implies that the conclusion of Theorem 2.1 of Bhattacharya and Patrangenaru (2005) regarding the CLT for sample Fréchet means holds, and we state it next.

Rather than working with the variables X_i , it is convenient to work with the “centered” variables $X_i^I = M^{-1/2} X_i M^{-1/2}$. By equivariance of the canonical geometric mean and average with respect to the group action of $GL(p)$ (Section 3.1) with $G = M^{-1/2}$, the population canonical geometric mean of the X_i^I 's is the identity matrix I (hence the superscript), and the canonical geometric average is $\bar{X}^I = M^{-1/2} \bar{X} M^{-1/2}$. To apply the CLT result above to the

centered variables, let the distance ρ needed there be the geodesic distance d_C (19) and let the chart $\phi(\cdot)$ be the Riemannian logarithmic map $\text{Log}_I(\cdot) = \log(\cdot)$ (40) at I , which covers $\text{Sym}^+(p)$ by (1). For two tangent vectors $U, \Theta \in \text{Sym}(p)$ on the tangent space at I , the map ϕ induces the distance

$$\rho^\phi(U, \Theta) = \rho(\phi^{-1}(U), \phi^{-1}(\Theta)) = d_C(e^U, e^\Theta), \quad U, \Theta \in \text{Sym}(p) \quad (42)$$

Then the image measure Q^ϕ has Fréchet mean $M^\phi = \phi(M^I) = 0$ with respect to this distance. Similarly, $\hat{M}_C^\phi = \phi(\hat{X}^I)$ is the Fréchet mean of the image empirical measure $\hat{Q}_n^\phi = (1/n) \sum_{i=1}^n \delta(X_i^\phi)$, where

$$X_i^\phi = \phi(X_i^I) = \log(M^{-1/2} X_i M^{-1/2}) = Y_i$$

is assumed to have finite covariance. Further, from (42), define the half squared distance map

$$\Theta \rightarrow G(U; \Theta) = \frac{1}{2}(\rho^\phi)^2(U, \Theta) = \frac{1}{2}d_C^2(e^U, e^\Theta) \quad (43)$$

together with its gradient and Hessian at Θ

$$\Psi(U; \Theta) = \nabla_\Theta G(U; \Theta) \in \text{Sym}(p), \quad H(U; \Theta) = \partial_\Theta \text{vecd}[\Psi(U; \Theta)] \in GL(q). \quad (44)$$

Note the factor 1/2 in front of (43), missing in Bhattacharya and Patrangenaru (2005, p. 1229). Both the gradient and Hessian exist because the map (43) is at least twice differentiable. Then Theorem 2.1 of Bhattacharya and Patrangenaru (2005) states that

$$\sqrt{n}(\hat{M}_C^\phi - M_C^\phi) = \sqrt{n}[\log(M^{-1/2} \hat{X} M^{-1/2}) - 0] \Rightarrow N(0, K^{-1} C K^{-1}),$$

where

$$C = \text{Cov}\{\text{vecd}[\Psi(X_1^\phi; 0)]\} \in \text{Sym}^+(q), \quad K = E[H(X_1^\phi; 0)] \in GL(q), \quad (45)$$

provided that K is invertible. The required quantities C and K are computed in detail, next.

B.2.2 Computation of the covariance C

To compute the gradient in (44), we follow the method of Moakher (2005). Let $\Theta + tA$ with $t \in \mathbb{R}$ and $A \in \text{Sym}(p)$ be a straight line in the tangent space at I . From (43), define the scalar function

$$h(t) = G(U, \Theta + tA) = \frac{1}{2}d_C^2(e^U, e^{\Theta+tA}) = \frac{1}{2}\text{tr}[\log^2(e^{-U}e^{\Theta+tA})] \quad (46)$$

where the last expression was obtained using (39) and (19). Then the gradient in (44) is the unique tangent vector $\Psi(U; \Theta)$ such that

$$\left. \frac{dh(t)}{dt} \right|_{t=0} = \langle \Psi(U; \Theta), A \rangle_I = \text{tr}(\Psi(U; \Theta)A), \quad (47)$$

where the Riemannian inner product (36) on the tangent space at I reduces to the Frobenius inner product (35).

Using the chain rule formula (Moakher, 2005, Prop. 2.1)

$$\frac{d}{dt}\text{tr}[\log^2 X(t)] = 2\text{tr}\left[\log X(t)X^{-1}(t)\frac{d}{dt}X(t)\right]$$

in (46) with $X(t) = e^{-U} e^{\Theta+tA}$ gives that the derivative of $h(t)$ at 0 is

$$\left. \frac{dh(t)}{dt} \right|_{t=0} = \text{tr} \left[\log(e^{-U} e^{\Theta}) \cdot e^{-\Theta} \left. \frac{d}{dt} e^{\Theta+tA} \right|_{t=0} \right]. \quad (48)$$

The first term inside the brackets has a polynomial expansion given by Goldberg's formula (34). Keeping terms only up to first order in Θ ,

$$\begin{aligned} \log(e^{-U} e^{\Theta}) &= -U + \Theta + \frac{1}{2}(-U\Theta + \Theta U) + \frac{1}{12}(U^2\Theta + \Theta U^2) - \frac{1}{6}U\Theta U \\ &\quad - \frac{1}{720}(U^4\Theta + \Theta U^4) + \frac{1}{180}(U\Theta U^3 + U^3\Theta U) - \frac{1}{120}U^2\Theta U^2 + \dots \end{aligned}$$

Similarly, keeping terms only up to first order in Θ and t in the second term in the brackets in (48),

$$e^{-\Theta} \left. \frac{d}{dt} e^{\Theta+tA} \right|_{t=0} = (I - \Theta + \dots) \frac{d}{dt} \left[I + \Theta + tA + \frac{1}{2}(\Theta + tA)^2 + \dots \right]_{t=0} = A + \frac{1}{2}(A\Theta - \Theta A) + \dots$$

Replacing in (48),

$$\begin{aligned} \left. \frac{dh(t)}{dt} \right|_{t=0} &= \text{tr} \left\{ \left[-U + \Theta + \frac{1}{12}(U^2\Theta - 2U\Theta U + \Theta U^2) \right. \right. \\ &\quad \left. \left. + \frac{1}{720}(U^4\Theta - 4U^3\Theta U + 6U^2\Theta U^2 - 4U\Theta U^3 + \Theta U^4) + \dots \right] A \right\}. \end{aligned}$$

Therefore, comparing with (47), we have that the desired gradient is

$$\begin{aligned} \Psi(U; \Theta) &= -U + \Theta + \frac{1}{12}(U^2\Theta - 2U\Theta U + \Theta U^2) \\ &\quad + \frac{1}{720}(U^4\Theta - 4U^3\Theta U + 6U^2\Theta U^2 - 4U\Theta U^3 + \Theta U^4) + \dots \end{aligned} \quad (49)$$

which is indeed in $\text{Sym}(p)$. In particular, $\Psi(X_1^\phi; 0) = -X_1^\phi = -Y_1$, with covariance $C = \text{Cov}(\text{vecd}(Y_1)) = \Sigma$.

B.2.3 Computation of the expected Hessian K

To obtain the Hessian in (44), we first vectorize the gradient (49) as

$$\begin{aligned} \text{vecd}[\Psi(U; \Theta)] &= -\text{vecd}(U) + \text{vecd}(\Theta) \\ &\quad + D' \left[\frac{1}{12}(U \otimes I - I \otimes U)^2 + \frac{1}{720}(U \otimes I - I \otimes U)^4 + \dots \right] D \text{vecd}(\Theta) \end{aligned}$$

where we have used property (31) and the property that $\text{vec}(AXB) = (B' \otimes A)\text{vec}(X)$. Thus

$$H(U; \Theta) = \frac{\partial \text{vecd}[\Psi(U; \Theta)]}{\partial [\text{vecd}(\Theta)]'} = I + D' \left[\frac{1}{12}(U \otimes U)^2 + \frac{1}{720}(U \otimes U)^4 + \dots \right] D$$

Finally, by (45) with $X_1^\phi = Y_1$, the desired matrix K is as given in (24). Note that K is PD and thus invertible.

References

- Anderson, T. W. (2003). *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons Ltd., Hoboken, New Jersey, third edition.
- Arsigny, V., Fillard, P., Pennec, X., and Ayache, N. (2006). Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magn Reson Med*, 56(2):411–421.
- Bammer, R., Auer, M., Keeling, S. L., Agustin, M., Stables, L. A., Prokesch, R. W., Stollberger, R., Moseley, M. E., and Fazekas, F. (2002). Diffusion tensor imaging using single-shot SENSE-EPI. *Magn Reson Med*, 48:128–136.
- Basser, P. J. and Pierpaoli, C. (1996). Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J Magn Reson B*, 111:209–219.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Statist Soc B*, 57(1):289–300.
- Bhattacharya, A. and Bhattacharya, R. N. (2008). Statistics on Riemannian manifolds: asymptotic distribution and curvature. *Proceedings of the American Mathematical Society*, 33(3):1225–1259.
- Bhattacharya, R. N., Elligson, L., Liu, X., Patrangenaru, V., and Crane, M. (2012). Extrinsic analysis on manifolds is computationally faster than intrinsic analysis, with applications to quality control by machine vision. *Appl. Stochastic Models Bus. Ind.*, 28:222–235.
- Bhattacharya, R. N. and Patrangenaru, V. (2003). Large sample theory of intrinsic and extrinsic sample means on manifolds. I. *Ann Statist*, 31(1):1–29.
- Bhattacharya, R. N. and Patrangenaru, V. (2005). Large sample theory of intrinsic and extrinsic sample means on manifolds. II. *Ann Statist*, 33(3):1225–1259.
- Chiu, T. Y. M., Leonard, T., and Tsui, K. W. (1996). The matrix-logarithmic covariance model. *J Amer Statist Assoc*, 91(433):198–210.
- DoCarmo, M. P. (1992). *Riemannian Geometry*. Birkhauser, Boston.
- Dougherty, R. F., Ben-Shachar, M., Deutsch, G. K., Hernandez, A., Fox, G. R., and Wandell, B. A. (2007). Temporal-callosal pathway diffusivity predicts phonological skills in children. *PNAS*, 104(20):8556–8561.
- Dougherty, R. F., Ben-Shachar, M., Deutsch, G. K., Potanina, P., and Wandell, B. A. (2005). Occipital-callosal pathways in children: Validation and atlas development. *Ann N Y Acad Sci*, 1064:98–112.
- Dryden, I. L., Koloydenko, A., and Zhou, D. (2009). Non-Euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *Ann of Appl Statist*, 3(3):1102–1123.
- Elligson, L., Groisser, D., Osborne, D., Patrangenaru, V., and Schwartzman, A. (2013). Non-parametric bootstrap of sample means of positive definite matrices with an application to diffusion tensor imaging data analysis. *J Math Imaging Vis*, page (submitted).
- Fletcher, P. T. and Joshi, S. (2007). Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing*, 87(2):250–262.

- Forstner, W. and Moonen, B. (1999). A metric for covariance matrices. In und Volker S. Schwarze, F. K., editor, *Quo vadis geodesia...?*, *Festschrift for Erik W. Grafarend on the occasion of his 60th birthday*, number 1999.6 in Technical Reports Department of Geodesy and Geoinformatics, pages 113–128. Stuttgart University.
- Fréchet, M. (1948). Les éléments aléatoires de nature quelconque dans un espace distancié. *Ann Inst H Poincaré*, 10:215–310.
- Goldberg, K. (1956). The formal power series for $\log(e^x e^y)$. *Duke Mathematical Journal*, 23:13–21.
- Groisser, D. (2004). Newton’s method, zeroes of vector fields, and the Riemannian center of mass. *Adv. Appl. Math.*, 33:95–135.
- Gupta, A. K. and Nagar, D. K. (2000). *Matrix Variate Distributions*. Chapman & Hall/CRC, Boca Raton, Florida.
- Johnson, N. L. and Kotz, S. (1970). *Continuous Univariate Distributions-1*. John Wiley & Sons, New York.
- Joshi, A. A., Shattuck, D. W., Thompson, P. M., and Leahy, R. M. (2007). Surface-constrained volumetric brain registration using harmonic mappings. *IEEE-Transactions in Medical Imaging*, 26(12):1657–1669.
- Kobayashi, S. and Nomizu, K. (1996). *Fundamentals of Differential Geometry - Volume II*. John Wiley & Sons, Inc., New York.
- Lang, S. (1999). *Fundamentals of Differential Geometry*. Springer-Verlag, New York.
- Lawson, J. D. and Lim, Y. (2001). The geometric mean, matrices, metrics, and more. *The American Mathematical Monthly*, November:797–812.
- LeBihan, D., Mangin, J.-F., Poupon, C., Clark, C. A., Pappata, S., Molko, N., and Chabriet, H. (2001). Diffusion tensor imaging: Concepts and applications. *J Magn Reson Imaging*, 13:534–546.
- Lenglet, C., Rousson, M., Deriche, R., and Faugeras, O. (2006). Statistics on the manifold of multivariate normal distributions: Theory and application to diffusion tensor mri processing. *J of Math Imaging and Vision*, 25(3):423–444.
- Leonard, T. and Hsu, J. S. J. (1992). Bayesian inference for a covariance matrix. *Ann Statist*, 20(4):1669–1696.
- Leporé, N., Brun, C., Chou, Y.-Y., Chiang, M.-C., Dutton, R. A., Hayashi, K. M., Lueders, E., Lopez, O. L., Aizenstein, H. J., Toga, A. W., Becker, J. T., and Thompson, P. M. (2008). Generalized tensor-based morphometry of HIV/AIDS using multivariate statistics on deformation tensors. *IEEE-Transactions in Medical Imaging*, 27(1):129–141.
- Mehta, M. L. (1991). *Random Matrices*. Academic Press, San Diego, California, 2nd edition.
- Moakher, M. (2002). Means and averaging in the group of rotations. *SIAM Matrix Anal Appl*, 24(1):1–16.

- Moakher, M. (2005). A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Matrix Anal Appl*, 26(3):735–747.
- Muirhead, R. J. (1982). *Aspects of Multivariate Statistical Theory*. John Wiley & Sons Ltd., Hoboken, New Jersey.
- Newman, M. and Thompson, R. C. (1987). Numerical values of Goldberg’s coefficients in the series for $\log(e^x e^y)$. *Mathematics of computation*, 48(177):265–271.
- Osborne, D., Patrangenaru, V., Ellingson, L., Groisser, D., and Schwartzman, A. (2013). Non-parametric two-sample tests on homogeneous Riemannian manifolds, cholesky decompositions and diffusion tensor image analysis. *J Multivariate Anal*, 119:163–175.
- Pasternak, O., Sochen, N., and Basser, P. J. (2010). The effect of metric selection on the analysis of diffusion tensor MRI data. *Neuroimage*, 49(3):2190–2204.
- Penneec, X. (1999). Probabilities and statistics on riemannian manifolds: Basic tools for geometric measurements. In *Int. Workshop on Nonlinear Signal and Image Processing (NSIP’99)*.
- Schott, J. R. (2005). *Matrix Analysis for Statistics*. John Wiley & Sons Ltd., Hoboken, New Jersey, 2nd edition.
- Schwartzman, A. (2006). *Random ellipsoids and false discovery rates: statistics for diffusion tensor imaging data*. PhD thesis, Stanford University.
- Schwartzman, A., Dougherty, R. F., and Taylor, J. E. (2010). Group comparison of eigenvalues and eigenvectors of diffusion tensors. *J Amer Statist Assoc*, 105(490):588–599.
- Schwartzman, A., Mascarenhas, W., and Taylor, J. E. (2008). Inference for eigenvalues and eigenvectors of Gaussian symmetric matrices. *Ann Statist*, 36(6):2886–2919.
- Smith, S. (2005). Covariance, subspace, and intrinsic Cramér-Rao bounds. *IEEE Trans Signal Proc*, 53(5):1610–1630.
- Whitcher, B., Wisco, J. J., Hadjikhani, N., and Tuch, D. S. (2007). Statistical group comparison of diffusion tensors via multivariate hypothesis testing. *Magn Reson Med*, 57:1065–1074.
- Zhu, H., Chen, Y., Ibrahim, J. G., Li, Y., Hall, C., and Lin, W. (2009). Intrinsic regression models for positive-definite matrices with applications to diffusion tensor imaging. *J Amer Statist Assoc*, 104(487):1203–1212.